

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 211 867 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
05.06.2002 Bulletin 2002/23

(51) Int Cl.7: H04L 29/06, H04L 29/12

(21) Application number: 01304560.4

(22) Date of filing: 23.05.2001

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 27.11.2000 JP 2000359810

(71) Applicant: FUJITSU LIMITED
Kawasaki-shi, Kanagawa 211-8588 (JP)

(72) Inventors:

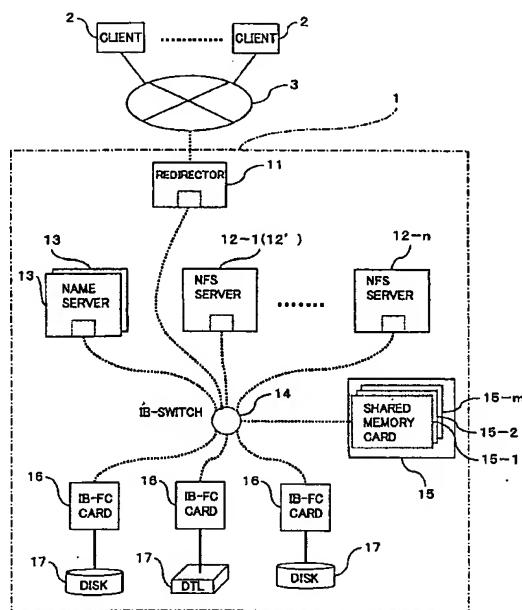
- Ooe, Kazuichi, c/o Fujitsu Limited
Kawasaki-shi, Kanagawa 211-8588 (JP)
- Nishikawa, Katsuhiko, c/o Fujitsu Limited
Kawasaki-shi, Kanagawa 211-8588 (JP)

(74) Representative: Hitching, Peter Matthew et al
Haseltine Lake & Co.,
Imperial House,
15-19 Kingsway
London WC2B 6UD (GB)

(54) Storage system

(57) There is provided a storage system including a storage unit (17) capable of storing therein file data, a plurality of file servers (12-1 to 12-n) for effecting file processing on the storage unit (17) in response to a file request concerning file data which is received from a client (2) through an external network (3), a file server administrating node (11) for unitarily administrating transfer processing for transferring the file data to the file server (12-i (i=1 to n)) based on the file request and reply processing for sending a reply message concerning the file request to the client (2) and an internal network (14) for interconnecting the storage unit (17), the file servers (12-i) and the file server administrating node (11) so that communication can be effected in the internal network (14). According to the foregoing arrangement, it is possible for the storage system to have satisfactory scalability capable of coping with the expansion of band of the network at a low cost.

FIG. 1



Description

[0001] The present invention relates to a storage system, and more particularly to a storage system which allows a plurality of clients involved in any network to share file data if the storage system is connected to the network.

[0002] As a conventional scheme of operating a network in which file data is shared by a plurality of nodes (clients) (hereinafter simply referred to as file sharing), there can be introduced a well-known manner in which, as for example schematically shown in FIG. 16, a file server 200 is built in any network 100 such as a LAN (Local Area Network) by utilizing an NFS (Network File System), and the file server 200 is connected with a secondary storage unit 400 through an interface 300 such as a so-called SCSI (Small Computer System Interface), and the file is shared by a plurality of clients 500 through the secondary storage unit 400.

[0003] The above-introduced manner, however, can encounter the following problems.

(1) Highly developed skill is requested for a person who is under duty of building the file server system and maintaining the same.

(2) It is not easy to expand the file server system (in its capacity, accessing performance or the like). If the system can be apparently expanded, one file server is unavoidably divided into a plurality of units and the entire number of units constituting the system becomes large, with the result that the maintenance cost for the system is also increased.

(3) Also, highly developed skill is requested for a person who is under duty of building the file server system and maintaining the same when the system suffers from any failure. Accordingly, the cost therefor will be expensive.

[0004] As a method for solving the problems, recently, there are proposed a method known as an NAS (Network Attached Storage). The NAS is equivalent to a unitarily built storage system including the file server 200 and the secondary storage unit 400 (see the portion surrounded by broken line in FIG. 16). If the NAS is connected to the network and simple setting operation is executed, the clients involved in the network can share data file, and further highly developed skill is unnecessary for a person under the duty of building the system and maintaining the same.

[0005] However, the above-described NAS still encounters a problem that it does not have satisfactory scalability for coping with the expansion of the transmission rate (in the current status, the transmission rate is of about 1Gbps, which is expected to reach about 10Gbps in several years in the future) of an LAN which is now progressively deployed. In other words, if the communication system is arranged so as to deal with the network as a connection destination which is in-

creasing in the transmission rate, it is unavoidably required to increase the number of file servers and secondary storage units provided within the NAS. As a result, a component which functions as the file server is divided into a plurality of components, and a component which functions as the secondary storage unit managed by the file server is also divided into a plurality of components.

[0006] That is, the above file server 200 and the secondary storage unit 400 are arranged as processing channels functioning in a parallel manner (independent manner). For this reason, it becomes necessary to carry out maintenance on each of the file server components, with the result that the maintenance cost will be increased.

[0007] Therefore, it is desirable to provide a storage system having satisfactory scalability capable of coping with the increase of transmission rate of the network with low cost.

[0008] According to the present invention there is provided a storage system including a storage unit capable of storing therein file data, a plurality of file servers for effecting file processing on the storage unit in response to a receive request, a file server administrating node for unitarily administrating transfer processing for transferring a request, received from the client via an external network, as the received request to the file server based on the request, and reply processing regarding the request to the client, and an internal network for interconnecting to the storage unit, the file servers, and the file server administrating node so that communication can be effected in the internal network.

[0009] According to the above storage system of the present invention, the network can be additionally provided with a file server or storage unit with ease depending on necessity. Furthermore, it is unnecessary to carry out maintenance on each of the file servers independently. Therefore, it becomes possible to provide satisfactory scalability in performance and capacity with low cost for an external network which transmission rate is increased.

[0010] In this case, the internal network may be connected with a name server which unitarily administrates the names of file data handled by the file servers. Also, the internal network may be connected with a shared memory which the file server administrating node and the file servers are made allowable to access. If the network is arranged to include the name server, in addition to the file server administrating node and the file servers, the name server will be made allowable to access the shared memory.

[0011] As described above, if the internal network is connected with the name server which can unitarily administrates names of all of the file data handled by the file servers, it is possible to create a name space which allows any of the file servers to have access to one and the same file via one and the same name path.

[0012] Furthermore, as described above, if the inter-

nal network is connected with the shared memory, the shared memory may be arranged to have handover information stored therein as required so that the file servers can exchange handover information through the shared memory to overcome a trouble which any of the file servers suffer from. With this arrangement, even if any of the file servers suffers from the trouble, the storage system will continue the normal operation. Therefore, the communication system becomes more durable against the problem in the network.

[0013] Further, according to the present invention, it is more preferable for the file server administrating node to be arranged to include a request analyzing unit for analyzing the contents of the request and a request transferring unit for transferring the request to a specified file server in accordance with the result of analysis of the request analyzing unit.

[0014] Further, the file server administrating node may be arranged so that a request regarding file data having an identical file data name is transferred to the same file server, based on data of the transferring operation history accumulated so far. According to the arrangement, processing speed will be remarkably improved. Further, the file server administrating node may be arranged to monitor the file servers to find a file server having relatively light load applied thereto and the request is transferred to the file server having relatively light load applied thereto. With this arrangement, each of the file servers will receive the requests evenly, thereby reliably avoiding a lot of load from being intensively applied to a particular file server, and avoiding a trouble caused by such an incident.

[0015] Further, file data, which request occurrence frequency is relatively high, may be cached in the main storage unit of the cache server so as to be processed by the cache server. With this arrangement, it becomes possible to remarkably reduce the access frequency to the storage unit, further improving the processing speed and processing performance.

[0016] In this case, the storage system may be arranged so that when the request occurrence frequency for the file data which is cached in the main storage unit of the cache server stays below a predetermined level, any server other than the cache server may succeed the processing of the request. With this arrangement, the cache server is relieved from holding the file data of which request occurrence frequency is no longer high for a long period of time. Therefore, the size of memory area, to be reserved in the cache server as the main storage unit, can be reduced, and the cache server will have more allowance in processing data, further improving the processing performance of the cache server.

[0017] Further, the file server administrating node may be arranged so that a header offset value indicative of the position of the boundary between the header portion and the substantial file data portion is calculated, and the header offset value is added to the request and transferred to the file server together with the request.

With this arrangement, in a network driver of the file server, the header portion and the data portion of the request can be copied on respective different regions for message which are handled by a higher rank layer

5 of the kernel, based on the header offset value. Therefore, no copy will be required in the kernel region, remarkably improving the processing speed and processing performance of the file server.

[0018] Further, the above-described file server administrating node may be arranged so that a reply message, to a request corresponding to a particular file data of which request occurrence frequency is relatively high, is cached. When a request corresponding to the file data is received, the cached reply message is returned to the client. With this arrangement, since the request need not be transferred to the file server, the responding speed for the client can be remarkably improved and hence the processing speed and the processing performance of the whole communication system can be dramatically improved.

[0019] The above-described storage unit may be arranged so that the storage unit permits access from an external node. With this arrangement, the storage system can be easily combined with another type of storage architecture. In this case, the name server may be also arranged to permit access from the external node. With this arrangement, the external node is allowed to have access to file data without arbitration control between the file server and the external node upon accessing the file.

[0020] Further, the above-described file server may be arranged to carry out file processing on the storage unit so as to respond to a request which is directly received from the external network. With this arrangement, the client is allowed to have access to the storage unit in a direct manner and in a manner via the file server. Also in this case, the storage system can be easily combined with another type of storage architecture.

[0021] Preferred features of the present invention will now be described, purely by way of example, with reference to the accompanying drawings, in which:-

FIG. 1 is a block diagram showing an arrangement of a storage system (storage architecture) as one embodiment of the present invention;
 FIG. 2 is a block diagram showing an arrangement of a redirector shown in FIG. 1;
 FIG. 3 is a block diagram for explaining an operation in which the redirector shown in FIG. 1 caches meta-information;
 FIG. 4 is a block diagram for explaining an operation in which the redirector shown in FIG. 1 returns a reply message;
 FIG. 5 is a block diagram showing an arrangement of an NFS server (name server) shown in FIG. 1;
 FIG. 6 is a block diagram for explaining an operation in which handover information of the NFS server is stored in a shared memory shown in FIG. 1 as a

backup;

FIG. 7 is a block diagram for explaining an operation in which an NFS server succeeds a task which was to be done by a down NFS on the bases of the handover information stored in the shared memory shown in FIG. 6 as a backup;

FIG. 8 is a block diagram for explaining an operation taken place when handover information of the name server is stored in the shared memory shown in FIG. 1 as a backup;

FIG. 9 is a block diagram for explaining an operation in which the redirector shown in FIG. 1 determines boundary information of a file access request;

FIG. 10 is a block diagram for explaining an operation in which the NFS server shown in FIG. 1 creates a zero copy status within the kernel, based on the boundary information of the file access request illustrated in FIG. 9;

FIG. 11 is a block diagram for explaining an operation in which the NFS server, shown in FIG. 1, creates a zero copy status within the kernel, based on the boundary information of the file access request illustrated in FIG. 9;

FIG. 12 is a diagram showing an example of format of the file access request shown in FIGS. 9 to 11; FIG. 13 is a block diagram showing an arrangement of the storage system shown in FIG. 1 in which an external node is allowed to have access to a secondary storage unit;

FIG. 14 is a block diagram showing an arrangement of the storage system, shown in FIG. 1, in which the external node is allowed to have access to the name server;

FIG. 15 is a block diagram showing an arrangement of the storage system, shown in FIG. 1, in which the external node is allowed to have access to the NFS server in a direct manner and in a manner by way of the redirector; and

FIG. 16 is a block diagram for explaining a conventional method of realizing file sharing among a plurality of nodes (clients) involved in a network.

(A) Description of First Embodiment

[0022] FIG. 1 is a block diagram showing an arrangement of a storage system (storage architecture) as one embodiment of the present invention. As shown in FIG. 1, a storage system 1 (hereinafter sometimes simply referred to as "system 1") is a system which allows a plurality of clients 2, connected to an external network 3 (e.g., gigabit Ethernet), to share file data. In order to attain this purpose, the storage system 1 includes a redirector 11, a plurality of NFS servers (file servers) 12-1 to 12-n, a name server 13, a shared memory 15, an IB-FC card 16, a secondary storage unit 17 (hard disk drive unit, tape recording unit (DTL) or the like) and so on. These components 11, 12-1, 13, 15, 16, and 17 are interconnected to one another through a high speed (internal)

network (Infiniband) switch 14 of which transfer speed is about 4 to 10Gbps, for example.

[0023] As described above, these components of the internal system (redirector 11, the NFS servers 12-i, the name server 13, the secondary storage unit 17 and so on) are connected to one another through the internal network 14 so that the network is composed of a number of clusters. Therefore, the network has a scalability for components such as the NFS server 12-i, the secondary storage unit 17 or the like, i.e., the NFS server 12-i, the secondary storage unit 17 or the like can be additionally connected to the internal network 14 with ease depending on necessity. Thus, the internal network comes to have a remarkably improved scalability in capacity, access performance or the like depending on the transmission speed of the external network 3 connected to the internal network.

[0024] The above-described redirector (file server administrating node) 11 is a unit for carrying out transfer processing such that various request messages (hereinafter simply referred to as "request"), received from an arbitrary client 2 through the external network 3, are transferred to any of the NFS servers 12-i (i = 1 to n). The redirector 11 also functions as a unit for unitarily administrating the reply message sending operation for each of the request to the clients 2 as a request source. That is, owing to the redirector 11, even if the NFS server 12-i, the secondary storage unit 17 or the like is additionally provided in the internal network as described above, it is unnecessary to rearrange the maintenance service mode for the additional components, unlike the conventional network.

[0025] The above term "request" means a request for file data (hereinafter sometimes simply referred to as "file") stored in the secondary storage unit 17. For example, the meaning of this term includes a request (file access request) for file operation to the substance of the file data and some other request such as a request for access to meta-information such as file name reference or the like. Further, each of the clients 2 is allowed to refer only to an IP (Internet Protocol) given to the redirector 11. That is, from each of the clients 2, the present system 1 can be observed as if the plurality of components constituting the system were a single integrated server.

[0026] In order to realize the above-mentioned functions, as for example shown in FIG. 2, the redirector 11 is arranged to include a gigabit Ethernet card 11a equipped with an interface for the external network 3, an infiniband card 11b equipped with an interface for the system internal components (internal network), a network processor 11c for intensively controlling the operation of respective cards 11a and 11b including the redirector 11 itself, and a memory (main storage unit) 11d for storing therein various data and a software (program) necessary for operating the network processor 11c.

[0027] The network processor 11c is interconnected to these components 11a to 11d through a PCI (Periph-

eral Component Interconnect) bus 11e so that they can communicate with one another.

[0028] In this case, the network processor 11c is a unit capable of receiving and sending (including conversion of protocol) a request and a reply message therefor exchanged between the internal network 14 and the external network 3. The network processor 11c is also capable of analyzing the request (protocol) received from the clients 2, deciding an access file name based on the analyzing result, and deciding a destination NFS server 12-i of the received request. According to the present embodiment, the network processor 11c is further designed to carry out the following control.

[0029] That is, the network processor 11c analyzes the request from the client 2, controls the request transfer operation so that jobs deriving from the requests are evenly distributed to respective servers 12-i (e.g., to distribute request preferentially to an NFS server 12-i having a relatively light load applied thereto), assigns a server to the request so that a request about the same file is allocated with the same server, preventing contention for file access in the NFS Servers 12-i.

[0030] To this end, the subject network processor 11c is arranged to include the following functional components as its main function.

- (1) A function as a request analyzing unit 11 for analyzing the contents of the request sent from the client 2.
- (2) A function as a request transferring unit 112 for transferring the received request to a specific NFS server 12-i in accordance with the result of analysis of the request analyzing unit 111.
- (3) A function as a transferring operation history recording unit 113 for recording transferring operation history of a request which was handled so far by the request transferring unit 112 (e.g., history of transferring the request to the memory 11d).
- (4) A function as a load monitoring unit 114 for periodically monitoring the load applying condition of respective NFS servers 12-i, by effecting an NFS server load monitoring daemon 115 as the background task.

[0031] With this arrangement, the request transferring unit 112 can transfer a request about a file having the same file name to the same NFS server 12-i based on the transferring operation history recorded in the transferring operation recording unit 113. Also, the request transferring unit 112 can transfer the received request to an NFS 12-i having relatively light load based on the result of load monitoring effected by the NFS server load monitoring daemon 115 (load monitoring unit 114).

[0032] Accordingly, the processing speed of the system will be remarkably improved. Furthermore, the requests can be evenly distributed to respective NFS servers 12-i. Therefore, each of the NFS servers 12-i can be reliably prevented from being applied with heavy load

intensively. Thus, the reliability of the system can be remarkably improved.

[0033] According to the arrangement of the subject network processor 11c, a cache memory 11f (memory 11d may replace for the cache memory 11f) is provided within the network processor 11c so that a handover message for access to meta-information is reserved in the cache memory 11f (see FIG. 3). When the redirector 11 receives a request of access to the meta-information from the client 2, the redirector 11 checks whether there is proper reply message stored in the cache memory 11f (or the memory 11d) or not. If the redirector 11 determines that there is a proper reply message stored in the cache memory 11f (or the memory 11d), the redirector 11 creates a reply message based on the reserved message and sends the reply message directly to the client 2 side (i.e., without sending the message through the NFS server 12-i or the name server 13) (see FIG. 4).

[0034] The above-described arrangement can be applied to not only the meta-information but file data. However, if the all kinds of file data are cached, the memory is required to have a capacity large enough to accommodate such storage. Therefore, it is preferred to arrange only file data, which has relatively high access occurrence frequency, is selected by the network processor 11c to be cached by the cache memory 11f (or the memory 11d).

[0035] That is, the cache memory 11f (or the memory 11d) functions as a cache unit for caching a reply message to be sent to the client about a particular file data which has relatively high request occurrence frequency. Further, the network processor 11c functions as a replying unit 116 (see FIG. 2) for returning the reply message stored in the cache memory 11f (or the memory 11d) to the client 2 when a new request is of the same data file which has been previously requested.

[0036] In this way, if the information (meta-information, file data or the like) is one having relatively high access occurrence frequency, a reply message for the information is cached on the side of the redirector 11 and the reply message is returned from the redirector 11 directly to the client 2 without being transferred to the NFS server 12-i. With this arrangement, the replying speed for the information, having relatively high access occurrence frequency to the client 2, can be remarkably improved, with the result that the processing speed and processing performance of the subject system 1 can be dramatically improved.

[0037] Now, arrangement of the NFS server 12-i according to the present embodiment will be hereinafter described. The NFS server 12-i is arranged to carry out file processing (e.g., writing, updating, reading and so on) in accordance the request sent from the redirector 11 by accessing the secondary storage unit 17 through the internal network (internal network switch) 14. The NFS server 12-i is also arranged to create a proper reply message and send the same to the redirector 11 so as to inform the file processing result to the client 2 which

is the request-source.

[0038] As for example shown in FIG. 5, the hardware arrangement of each of the NFS server 12-i is provided with a CPU (Central Processing Unit) 12a, a memory (main storage unit) 12b and an interface card (IB-IF) 12c equipped with an interface (protocol conversion) with the internal network 14. When the CPU 12a reads an NFS server software (program) stored in the memory 12b, the NFS server 12-i is operated so as to realize the foregoing function.

[0039] At this time, if respective servers 12-i manage the file name independently, file data can be attached with different management file names by respective NFS servers 12-i in spite of the fact that the file data has the same substance. Conversely, file data can be attached with the same management file name by respective NFS servers 12-i in spite of the fact that each of the file data has different substance, thereby leading to contention for file access among the NFS servers 12-i.

[0040] In order to avoid the above-described inconvenience, the name server 13 is introduced into the system. That is, the name server 13 places meta-information access from all of the NFS servers 12-i under the unified control, whereby the file data handled by all NFS servers 12-i are given file names based on the unified control. Thus, the file access contention among the NFS servers 12-i can be avoided. If the system is provided with the name server 13, the reliability of the subject system 1 can be remarkably improved in file sharing.

[0041] As shown in FIG. 1, the subject storage system 1 is provided with two name servers 13, or a name server for current use and a name server for spare use, in order that the storage system 1 can cope with any abnormal incident such as a trouble (down) or the like. Further, these name servers 13 are made to have the same arrangement as that of the NFS servers 12-i (see FIG. 5) in terms of hardware arrangement. That is, each of the name servers 13 is arranged to include a CPU 13a, a memory (main storage unit) 13b, and an interface card 13c equipped with an interface with the internal network 14. Also in this case, the name server software (program) stored in the memory 13b is read by the CPU 13a, and the CPU 13b is operated in accordance with the read program, whereby the above-described function of the name server 13 is realized.

[0042] Now, the arrangement of the shared memory 15 will be hereinafter described. The shared memory 15 is a memory unit which can be accessed from each of the redirector 11, the NFS server 12-i and the name server 13 through the internal network 14. For example, when a certain NFS server 12-i or the working name server 13 comes to be down (suffers from failure), the task which is being done in the NFS server 12-i or the name server 13 should be relayed to another NFS server 12-k ($k=1$ to n , $K \neq i$) or the spare server 13. The shared memory 15 has information necessary for the NFS server 12-i or the name server 13 to relay the task stored therein so that each of the information pieces is held (as

a backup) in a memory card (shared memory card) 15-1 to 15-m (m is a natural number) for the server 12-i and the name server 13 independently (see FIGS. 6 and 8).

[0043] That is, the above-described NFS server 12-i (CPU 12a) or the name server 13 (CPU 13a) is arranged to include a handover information recording unit 121 (131) (see FIG. 5) for recording, in the shared memory 15, information which is necessary for the NFS server 12-i or the name server 13 to relay the task to be done to another NFS server 12-i or the spare name server 13 in order to cope with any abnormal incident brought about in the network.

[0044] As for example schematically shown in FIG. 6, the trouble causing that the NFS server 12-i is down is detected by effecting the monitoring of the NFS server monitoring daemon 132 on the background by the working name server 13 (CPU 13a). On the other hand, the trouble causing that the working name server 13 is down is detected by effecting the monitoring of the NFS server monitoring daemon 133 on the background by the spare name server 13 (CPU 13a).

[0045] As schematically shown in FIG. 7, when it is detected that the NFS server 12-i comes to be down (Step S1), the working name server 13 (CPU 13a) issues an instruction to an NFS server 12-k other than the down NFS server 12-i (e.g., an NFS server 12-k having relatively light load applied thereto) so that the NFS server 12-k succeeds the task which is being done by the down NFS server 12-i. Also, the working name server 13 (CPU 13a) notifies the redirector 11 that the NFS server 12-i is down (Step 2).

[0046] In this way, upon receiving the instruction of succeeding the task, the NFS server 12-k (CPU 12a) has access to the shared memory 15 through the internal network 14, and reads handover information stored in the shared memory 15 as backup, and then succeeds the task of the NFS server 12-i which was down due to trouble (Step S3). At this time, the redirector 11 (network processor 11c) receives the above notification from the name server 13, and in response to the receiving of the notification, the redirector 11 is prohibited from transferring a request, by the request transferring unit 112, to the down NFS server 12-i.

[0047] Namely, in this case, the name server 13 (CPU 13a) is arranged to include functions of an abnormal incident detecting unit 134 (see FIG. 6) for detecting an abnormal incident brought about in the NFS server 12-i, and as a relaying instruction generating unit 135 (see FIG. 6) for generating an instruction to the NFS server 12-k other than the NFS server 12-i so that the NFS server 12-k succeeds the task of the NFS server 12-i suffering from the abnormal incident based on the handover information stored in the shared memory 15, when the abnormal incident detecting unit 134 detects the abnormal incident in the NFS server 12-i.

[0048] As described above, according to the arrangement of the present storage system, even if one of the NFS server 12-i or the name server 13 is down, the NFS

server 12-k other than the NFS server 12-i or the name sever 13 for the spare use can succeed the task being done in the NFS server 12-i or the name server 13. Therefore, the storage system 1, as a whole, can continue the file processing in a normal manner. Accordingly, the anti-trouble resistance of the storage system can be remarkably improved.

[0049] Although descriptions have been made on the redundancy of the NFS servers 12-i and the name sever 13, it is needless to say that the redirector 11 may be similarly made redundant. Furthermore, while in the above described arrangement the spare use name server 13 succeeds the current name server 13 when the current name server 13 is down, any of the NFS server 12-i may succeed the task of the current name server 13.

[0050] Now, description will be concretely made on the processing carried out in the redirector 11 and the NFS server 12-i when the request is transferred from the redirector 11 to the NFS server 12-i.

[0051] When the redirector 11 receives from the client 2, a file access request indicating that file data is to be written, for example, the redirector 11 analyzes the file access request with the request analyzing unit 111.

[0052] As for example shown in FIG. 12, the above-described file access request is made up of a header portion 21 and a substantial file data portion 22, the header portion 21 being composed of a physical layer header (Phy Header) 21a, an IP header (Internet Protocol Header) 21b, a TCP header (Transmission Control Protocol Header) 21c, an NSF header 21d and so on, the substantial file data portion 22 being composed of substantial file data to be written into the secondary storage unit 17.

[0053] As schematically shown in FIG. 9, the request analyzing unit 111 determines a position at which the substantial file data in the file access request starts, i.e., the boundary between the header portion 21 and the substantial file data portion 22 as a header offset value [boundary information; e.g., the number of bits "a" counted from the head portion of the file access request message] 23. The thus determined boundary information 23 is sent to the request transferring unit 112, and the request transferring unit 112 sends the file access request together with the boundary information 23 attached thereto, to an NFS server 12-i as a transfer destination.

[0054] That is, as shown in FIG. 2, the request analyzing unit 111 is arranged to include functions of a header offset value analyzing unit 111a and a header offset value adding unit 111b, the header offset value analyzing unit 111a analyzing the received file access request and determining the header offset value 23 indicative of the boundary position between the header portion 21 and the substantial file data portion 22 of the file access request, the header offset value adding unit 111b attaching the header offset value 23 obtained by the header offset value analyzing unit 111a to the file access re-

quest which is to be sent to the NFS server 12-i.

[0055] Thereafter, on the side of the NFS server 12-i, the NIC (Network Interface Card) driver (network driver) 122 allocates starting addresses of the substantial file data portion 22 and regions other than the substantial file data portion 22 to page boundaries (page boundary (separated region); buffers (mbuf) 123 and 124) which are handled by within-kernel higher rank layer (kernel higher rank layer) (NFS processing layer) based on the header offset value 23 added to the file access request on the side of the redirector 11 as described above (see FIG. 10).

[0056] According to the above-described processing scheme, as for example schematically shown in FIG. 11, when the file access request reaches the file system unit 125 in the kernel higher rank layer, the starting address (pointer) of the substantial file data portion 22 is brought to the pointer for replacing the pointer which is guiding the data to the file system buffer 126. Only with this operation, the data can be transferred to the file system buffer 126 without copying (map switching) the data. In other words, no copy can be achieved within the kernel area. Accordingly, DMA (Direct Memory Access) can be realized at a high speed, and the processing speed and processing performance of the NFS server 12-i can be remarkably improved.

[0057] In the above case, it is considered that the boundary between the header portion 21 and the substantial file data portion 22 is determined on the side of the NIC driver 122. In this case, however, the amount of task (header analyzing) imposed on the NIC driver 122 will be increased. (Usually, the NIC driver 122 analyzes only the physical layer header 21a.) Therefore, as described above, it is recommended that the redirector 11 is arranged to have the analyzing function (request analyzing unit 111) for analyzing the header portion 21 from the first, and the boundary is determined on the side of the redirector 11. With this arrangement, the layer of the NIC driver can be protected from a heavy task (and hence the layer of the NIC driver can be free from being lowered in its processing capability) and kernel zero copy can be achieved within the higher rank layer (NFS processing layer).

[0058] According to the arrangement of the present embodiment as described above, the storage system 1 is arranged to include within the system 1, the redirector 11, the plurality of NFS servers 12-i, the name server 13, the shared memory 15, and the secondary storage unit 17, and these components are connected by means of the high speed internal network 14. Therefore, the component such as the NFS server 12-i or the secondary storage unit 17 can be additionally provided in the network depending on necessity without any difficulties. Furthermore, since each of the NFS servers 12-i does not require independent maintenance operation, the storage system 1 can secure scalability in terms of performance and capacity which can cope with the enlargement of band of the external network 3 (e.g., a local area

network operating at a transmission rate of up to 10Gbps with sufficiently low cost.

[0059] In particular, according to the above-described arrangement, the redirector 11 controls the task delivery scheme. That is, the amount of task is delivered to respective NFS servers 12-i so that each of the NFS servers 12-i is supplied with substantially an equal amount of task, the task of a request concerning the same file data is assigned to the same NFS server 12-i, and a specially cached reply message is prepared not on the side of NFS server 12-i but on the side of redirector 11 so as to respond to a request of a file of which access occurrence frequency is relatively high. With this arrangement, the processing speed and processing performance of the system is dramatically improved, thereby realizing scalability in terms of performance and capacity which can cope with a LAN operating at a transmission rate of up to 10Gbps

(B) Description of First Modification

[0060] While in the above-described storage system 1 the memory 12b provided in each of the NFS servers 12-i is not particularly arranged in terms of its capacity, one of the memory 12b may be arranged to have a capacity larger than that of the memory of any other NFS servers 12-i so that the NFS server is made to function as a cache server 12' (see FIG. 1), and that fundamentally any file access by way of the cache server 12' is processed by only writing and reading on the memory 12b so as to respond to the side of the client 2.

[0061] If it is detected that a certain file request becomes one having access occurrence frequency in a predetermined period of time is higher than a certain level (threshold level), a special reply message is cached in the memory 12b so that the cache server 12' responds to the request.

[0062] This scheme will be more concretely described. Initially, the access occurrence frequency of respective files is monitored on the side of the redirector 11. If any file having access occurrence frequency higher than a certain threshold level is found, the redirector 11 issues an instruction to the name server 13, the NFS servers 12-i and the cache server 12' so that the access concerning the file having access occurrence frequency higher than the certain threshold level is to be processed in the cache server 12'.

[0063] At this time, the redirector 11 (request transferring unit 112) is arranged to transfer a request concerning a file of which access occurrence frequency is relatively high, to the cache server 12'. With this arrangement, the cache server 12' is exclusively obliged to carry out the task of the request concerning a file of which access occurrence frequency is relatively high, without accessing to the secondary storage unit 17. Therefore, great contribution can be expected to the improvement in the processing speed and processing performance of the storage system 1.

[0064] The task delivery scheme of the redirector 11 may be further arranged as described below. That is, if the access occurrence frequency of a certain file is lowered (i.e., the access occurrence frequency of the file

5 cached in the memory 12b of the cache server 12' falls below a certain threshold level), the redirector 11 issues an instruction to the name server 13, the NFS servers 12-i and the cache server 12' so that a proper NFS server (e.g., a server having relatively light task applied thereto) is designated and the task of processing the request concerning the certain file is assigned to the designated server.

[0065] In this case, the redirector 11 (request transferring unit 112) changes the transferring destination of 15 the request from the cache server 12' to any servers 12-i other than the cache server 12'. Thus, the file data of which access occurrence frequency is lowering can be prevented from being cached in the cache server 12'. As a result, the memory capacity necessary for the 20 cache server 12' can be reduced. Alternatively, the cache server 12' comes to have more allowance in processing tasks. Therefore, the processing performance can be more improved.

(C) Description of Second Modification

[0066] The connection arrangement of the secondary storage unit 17 may be further modified as, for example, shown in FIG. 13. That is, the secondary storage unit 30 17 is connected to the name server 13 and the NFS server 12-i by way of an FC switch 18, thereby to form a secondary storage unit network, and the FC switch 18 is connected to the external node 19. Thus, the secondary storage unit 17 permits access from the external 35 node 19. In this case, however, the file system operated by the external node 19 is requested to be the same as the file system within the storage system 1 (in the above-described case, the frame portions applied with halftone notation in FIG. 13 are the objects of the requirement).

[0067] If the connection arrangement around the secondary storage unit 17 is modified as described above, 40 in order to avoid contention in access with the NFS servers 12-i provided within the system 1, any arbitration control is requested on the access from the external node 19. However, external node 19 is allowed to have 45 access to the file provided in the storage system 1.

[0068] However, as shown in FIG. 14, for example, if the system is arranged so that the external node 19 is allowed to have access to the name server 13, the access 50 from the external node 19 will be placed under the control of file name management effected in the system 1. Therefore, the external node 19 is allowed to have access to the files provided in the system 1 without the above-described arbitration control. While the above 55 description is made with reference to FIG. 14 on the case in which the name server 13 is allowed to have access by way of the internal network 14, it is needless to say that, as shown in FIG. 13, the name server 13

may be allowed to have access by way of the secondary storage unit (FC switch 18).

[0069] While in the above arrangement, the NFS servers 12-i and the external node 19 are separately provided, the NFS servers 12-i and the external node 19 may be unitarily arranged as, for example, shown in FIG. 15. In other words, the file server 12-i is arranged to effect file processing, corresponding to the request is directly received from the external network 3, on the secondary storage unit.

[0070] According to the above arrangement, if a certain client 2 tries to access the storage system 1 by way of the redirector 11, the NFS server 12-i functions as the above-described file server of the storage system 1. If the client 2 tries to access the storage system 1 directly, the NFS server 12-i functions as an ordinary server which responds to the client not by way of the redirector 11 but directly. In other words, the above-described arrangement allows the clients 2 to access in two manners, i.e., to access by way of the NFS server 12-i and to access not by way of the NFS server 12-i.

[0071] In either of the cases, the present system allows the direct access from the outside. Therefore, it becomes possible to realize harmonization with another storage architecture (e.g., SAN (Storage Area Network) or the like). In FIGS. 14 and 15, reference numeral 20 represents a network disk adapter equipped with an interface component providing interface between the internal network 14 and the secondary storage unit 17.

[0072] While in the arrangements illustrated in FIGS. 13 to 15 the aforementioned shared memory 15 is omitted, it is needless to say that the shared memory 15 may be provided in the arrangements. If these arrangements are provided with the shared memory 15, backup processing can be effected even in the arrangement illustrated in FIGS. 13 to 15.

(D) Other Disclosure

[0073] While in the above embodiment descriptions have been made on a case in which the infiniband is applied as the internal network 4 and gigabit Ethernet is applied as the external network 3, it is needless to say that any high speed network other than these types of network can be employed for system building.

[0074] Further, the name server 13 or the shared memory 15 should not always be provided. That is, the purpose of the present invention can be satisfactorily achieved even if either of or both of these components are omitted. Furthermore, while in the above embodiments an NFS is employed as the file server, the present invention is not limited to this arrangement but other type of file system can be employed without any difficulties.

[0075] While in the above embodiment it is presupposed that the internal network 14 has a capacity (transmission rate) of about 4 to 10Gbps, the internal network 14 can cope with the expansion of band of the external network 3 by variably settling the transmission rate de-

pending on the expansion of band of the external network 3.

[0076] While several embodiments and modifications have been described above, the present invention is not limited to these embodiments but various changes and modifications other than the above embodiments can be effected without departing from the scope of the present invention.

10 Claims

1. A storage system comprising:

15 a storage unit (17) capable of storing file data; a plurality of file servers (12-1 to 12-n) for effecting file processing on the storage unit (17) in response to a received request; a file server administrating node (11) for unitarily administrating transfer processing for transferring a request, received from the client (2) via an external network (3), as said received request to the file servers (12-i: i=1 to n), and reply processing for sending a reply message regarding the request to the client (2); and an internal network (14) for interconnecting the storage unit (17), the file servers (12-i), and the file server administrating node (11) so that communication can be effected in the internal network (14).

20 2. A storage system according to claim 1, characterized in that the internal network (14) is connected with a name server (13) for unitarily administrating the names of file data handled by the file servers (12-i).

25 3. A storage system according to claim 2, characterized in that the internal network (14) is connected with a shared memory (15) which can be accessed from the file server administrating node (11) and the file servers (12-i).

30 4. A storage system according to claim 2, wherein the internal network (14) is connected with a shared memory (15) which can be accessed from the file server administrating node (11), the file servers (12-i) and the name server (13).

35 5. A storage system according to any of claims 1 to 4, characterized in that the file server administrating node (11) comprises a request analyzing unit (111) for analyzing the contents of the request and a request transferring unit (112) for transferring the request to a specified file server (12-i) in accordance with the result of analysis obtained by the request analyzing unit (111).

6. A storage system according to claim 5, characterized in that the file server administrating node (11) comprises a transferring operation history recording unit (113) for recording transferring operation history of a request which was previously handled by the request transferring unit (112), and
 the request transferring unit (112) is arranged to transfer a request regarding file data having an identical file data name to the same file server (12-i), based on the transferring operation history recorded by the transferring operation history recording unit (113).
 5
7. A storage system according to claim 5, characterized in that the file server administrating node (11) comprises a load monitoring unit (114) for monitoring load applied on the file server (12-i), and that
 the request transferring unit (112) is arranged to transfer the request to a file server (12-i) having relatively light load applied thereto, based on the result of monitoring monitored by the load monitoring unit (114).
 10
8. A storage system according to claim 5, characterized in that at least one of the file servers (12-i) comprises a main storage unit (12b) capable of caching file data of the storage unit (17), whereby the file server (12-i) can function as a cache server (12') for executing file processing in the main storage unit (12b) in accordance with the request.
 15
9. A storage system according to claim 8, characterized in that the main storage unit (12b) of the cache server (12') is arranged to cache file data of which request occurrence frequency within a constant time period is equal to or more than a predetermined level, and that
 the request transferring unit (112) is arranged to transfer the request regarding the file data of which request occurrence frequency is equal to or more than a predetermined level, to the cache server (12').
 20
10. A storage system according to claim 9, characterized in that if it is detected that the request occurrence frequency of the file data cached in the main storage unit (12b) of the cache server (12') becomes equal to or less than a predetermined level, the request transferring unit (112) changes the destination of the request from the cache server (12') to any file server (12-i) other than the cache server (12').
 25
11. A storage system according to claim 5, characterized in that the request analyzing unit (111) comprises a header offset value analyzing unit (111a) for analyzing the request so as to calculate a header offset value indicative of the position of the bound-
 30
- ary between the header portion (21) and the substantial file data portion (22), and
 a header offset value adding unit (111b) for adding the header offset value obtained by the header offset value analyzing unit (111a) to the request which is transferred to the file server (12-i).
 35
12. A storage system according to claim 11, characterized in that the file server (12-i) comprises a network driver (122) for copying the header portion (21) and the substantial file data portion (22) of the request to respective different regions for message which are handled by a higher rank layer of the kernel.
 40
13. A storage system according to any of claims 1 to 12, characterized in that the file server administrating node (11) comprises a caching unit (11f) for caching a reply message for responding to the client (2) who sends a request concerning a particular file data having relatively high request occurrence frequency, and that
 a responding unit (116) for responding to the client (2) with the reply message stored in the caching unit (11f) if it is determined that the request is one concerning the particular file data.
 45
14. A storage system according to claim 3 or 4, characterized in that the file server (12-i) is provided with a handover information recording unit (121) for recording handover information, which is necessary for another file server (12-j;j=1 to n) to succeed the task of that file server (12-i) in preparation for an abnormal incident, in the shared memory (15).
 50
15. A storage system according to claim 14, characterized by further comprising:
 55
- an abnormal incident detecting unit (121) for detecting abnormal incident occurred in the file server (12-i); and
 a handover instructing unit (135) for issuing instruction to a file server (12-j) other than that file server (12-i) (hereinafter referred to as abnormal file server) so that the file server (12-j) other than that file server (12-i) succeeds the task of the abnormal file server (12-i) based on the handover information stored in the shared memory (15), when the abnormal incident detecting unit (121) detects an abnormal incident which has occurred in the file server (12-i).
 60
16. A storage system according to any of claims 1 to 15, characterized in that the storage unit (17) is arranged to allow the access from the external node (19).
 65
17. A storage system according to any of claims 2 to

15, characterized in that the name server (13) is arranged to allow the access from the external node (19).

18. A storage system according to any of claims 1 to 5
17, characterized in that the file server (12-i) is arranged to effect file processing on the storage unit (17) if a request corresponding to the file processing is directly received from the external network (3).

10

15

20

25

30

35

40

45

50

55

11

FIG. 1

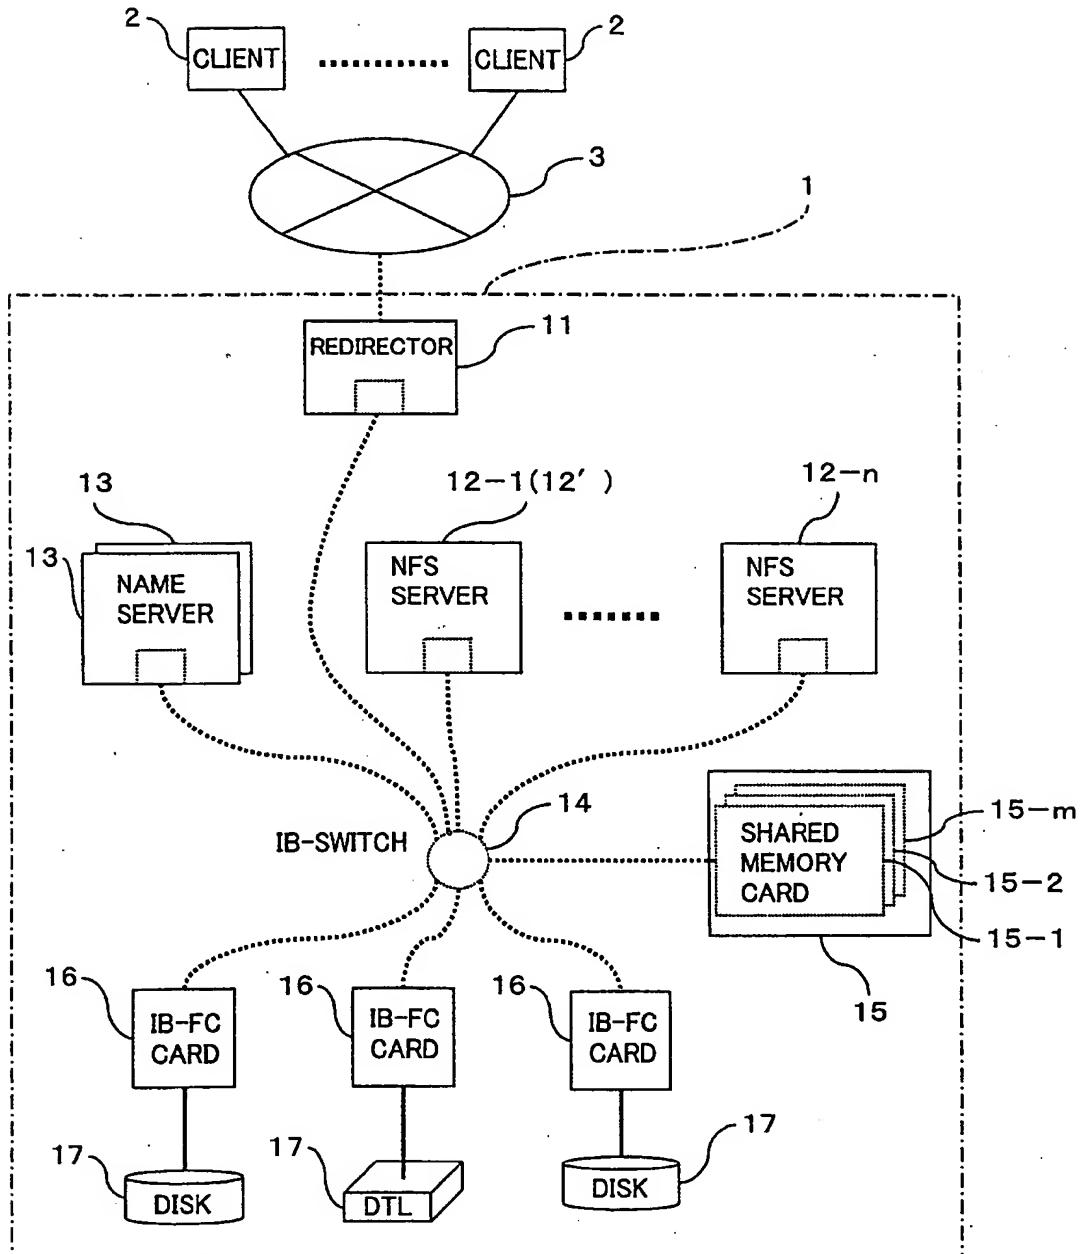


FIG. 2

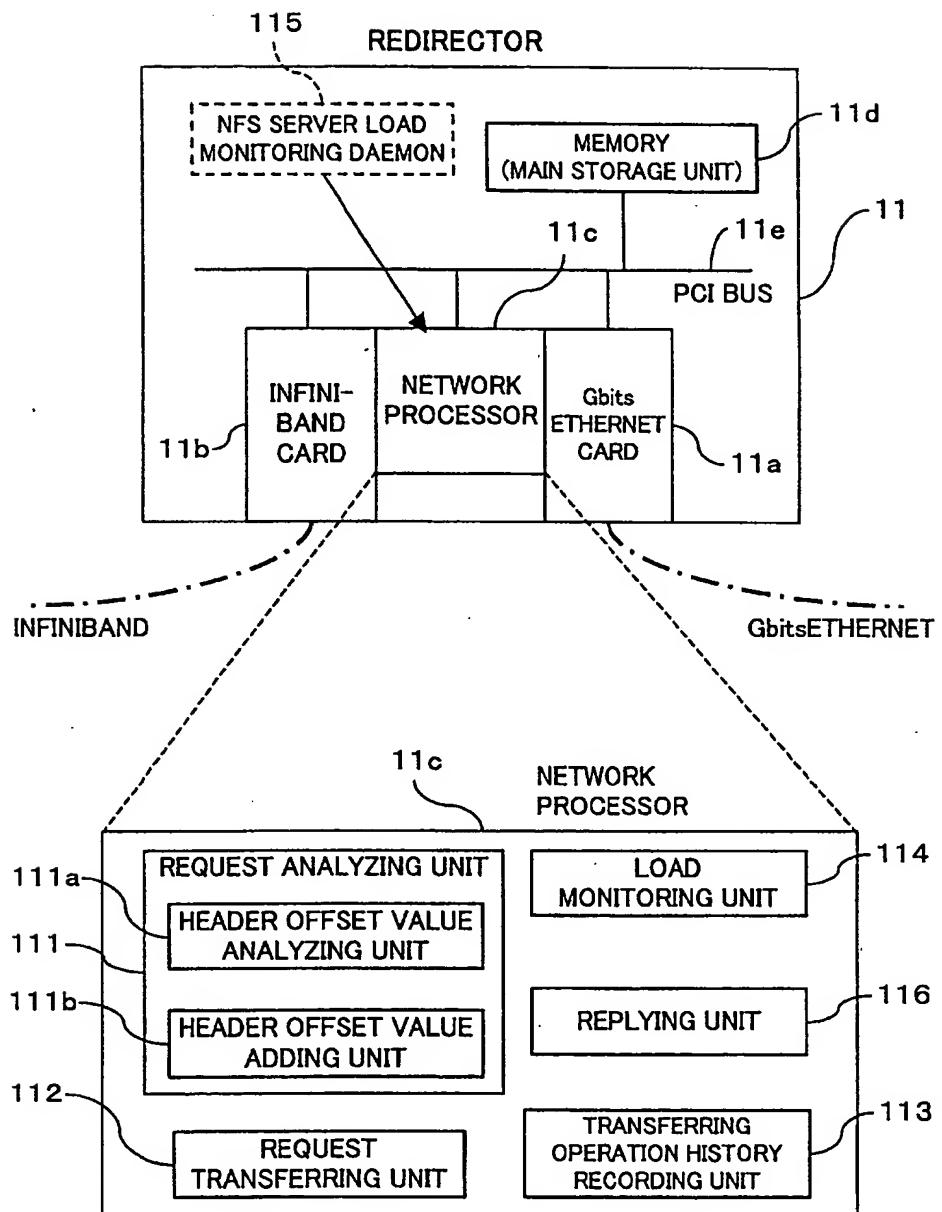


FIG. 3

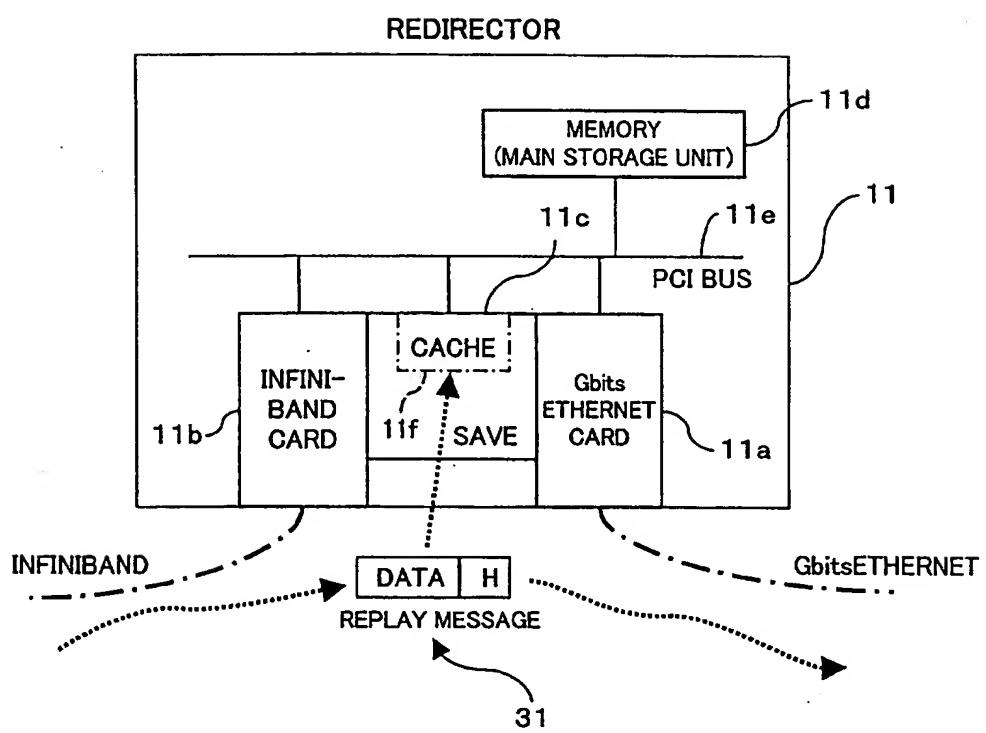


FIG. 4

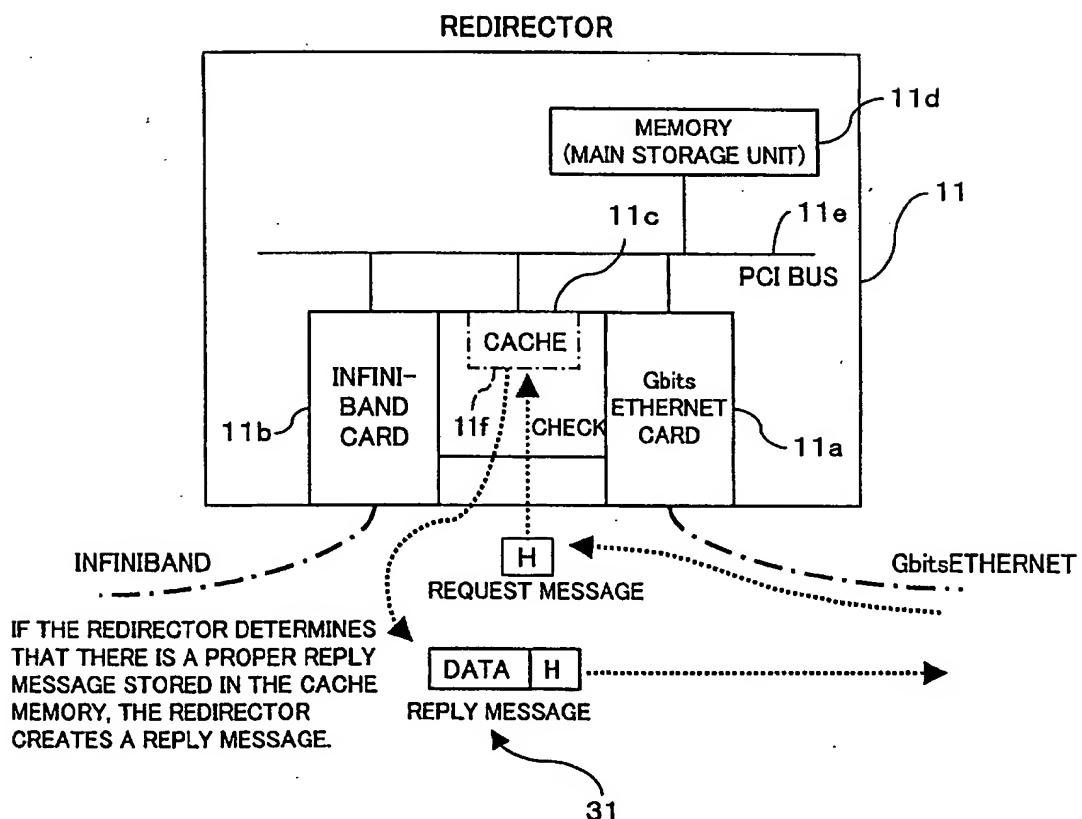
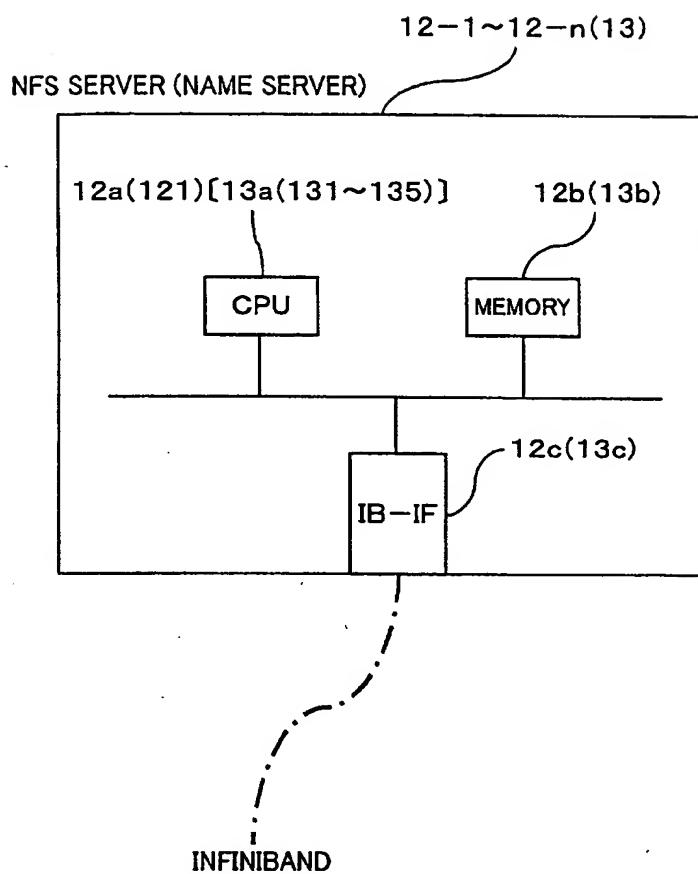


FIG. 5



121, 131:HANOVER INFORMATION RECORDING UNIT
132:NFS SERVER MONITORING DAEMON
133:NAME SERVER MONITORING DAEMON
134:ABNORMAL INCIDENT DETECTING UNIT
135:HANOVER INSTRUCTING UNIT

FIG. 6

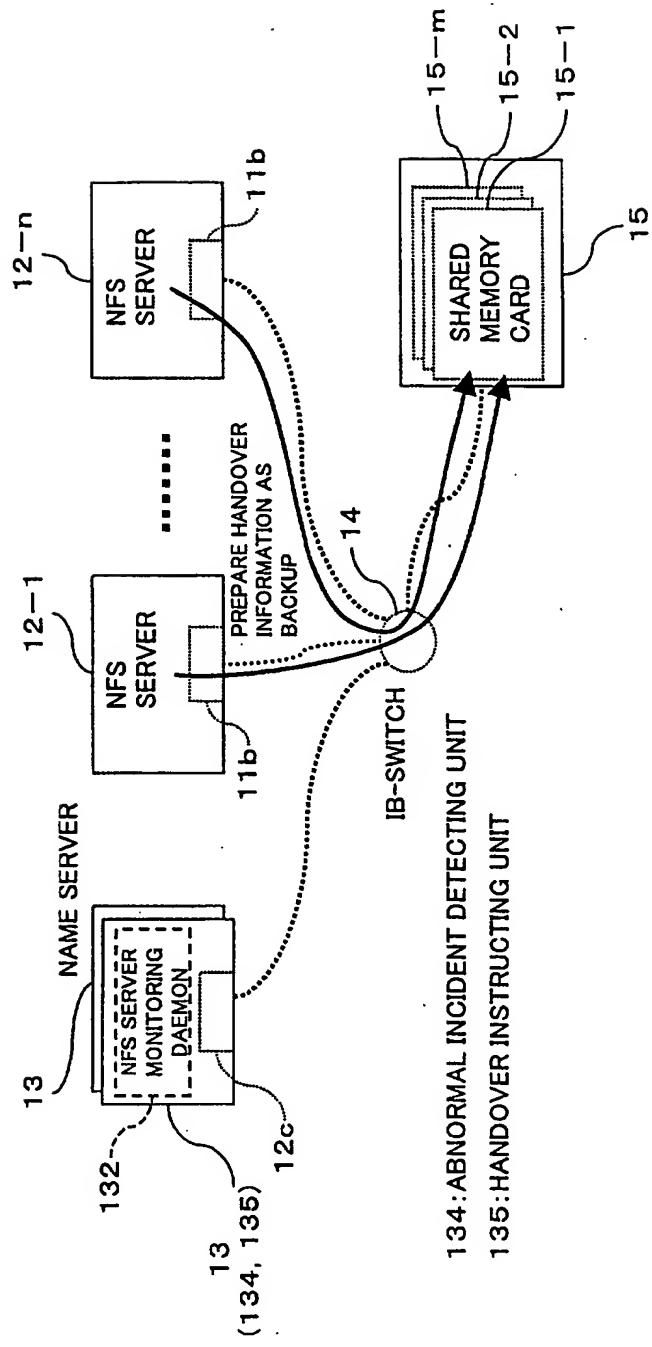


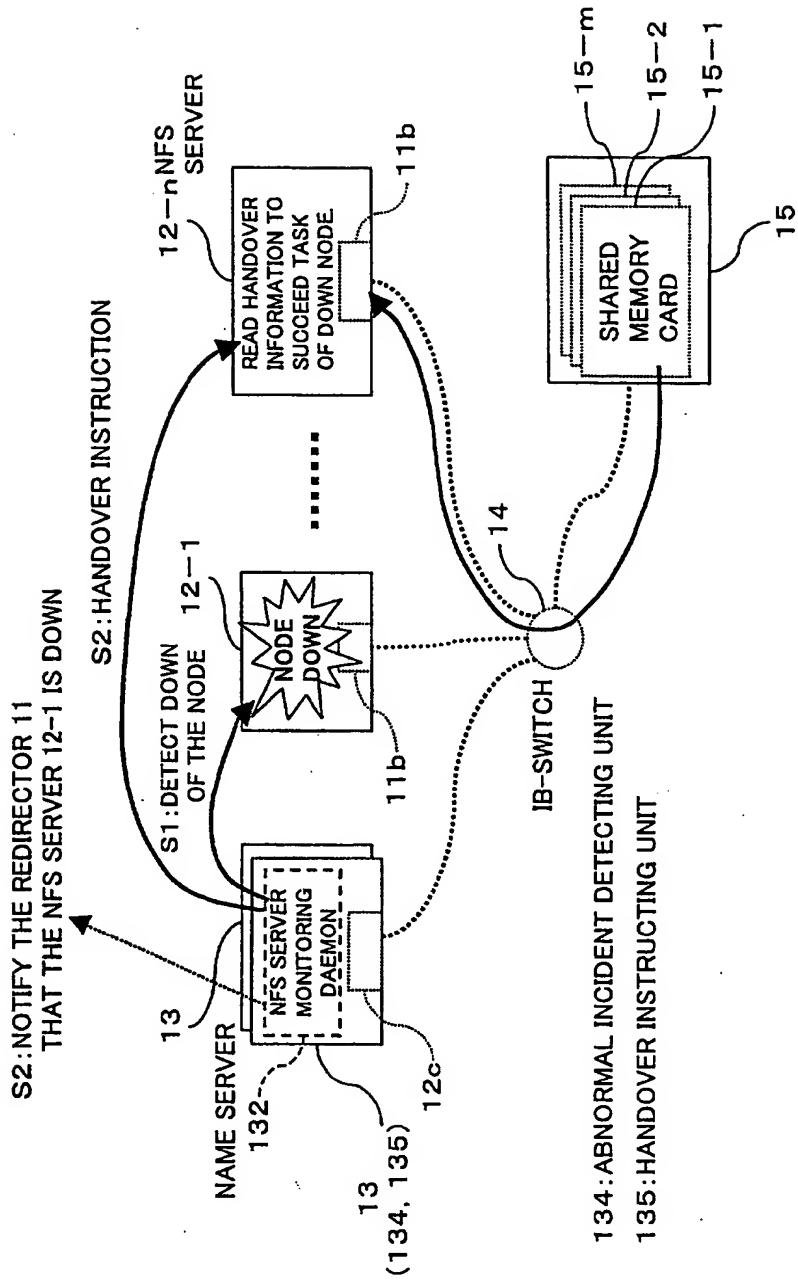
FIG. 7

FIG. 8

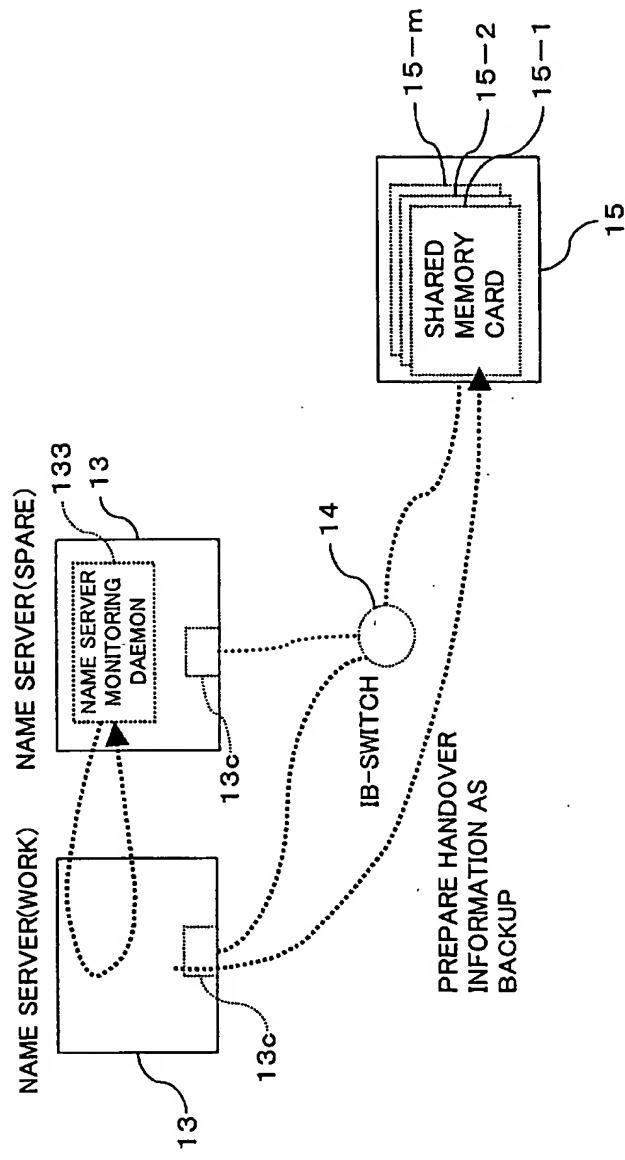


FIG. 9

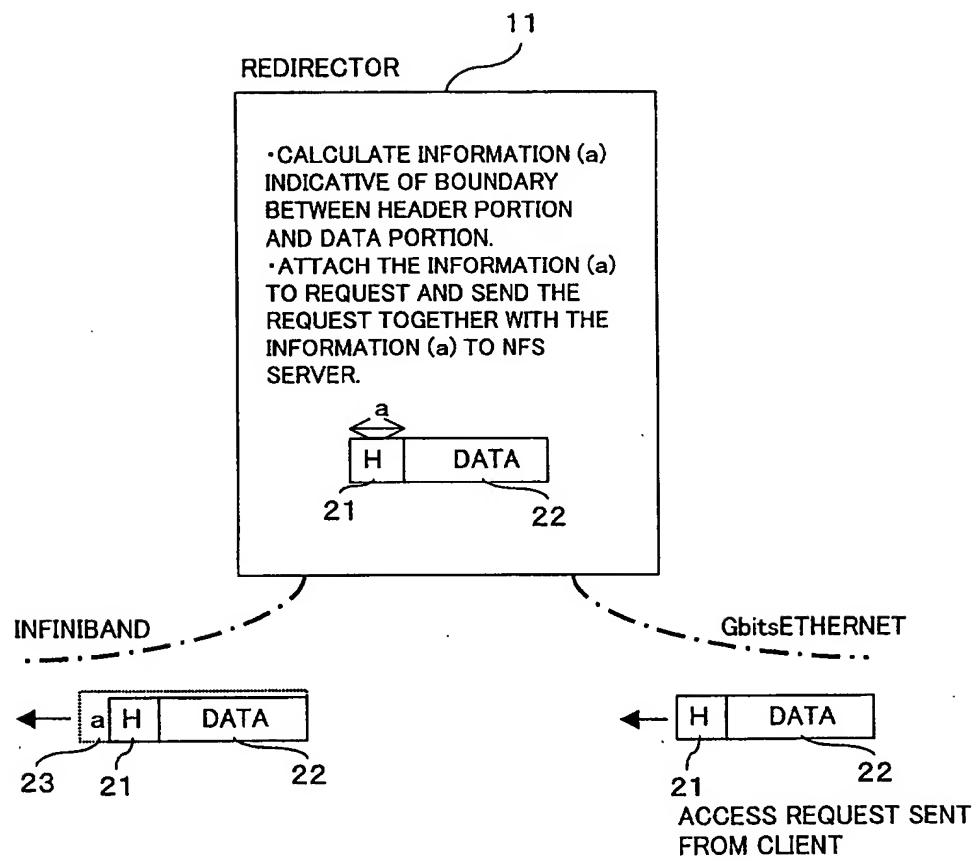


FIG. 10

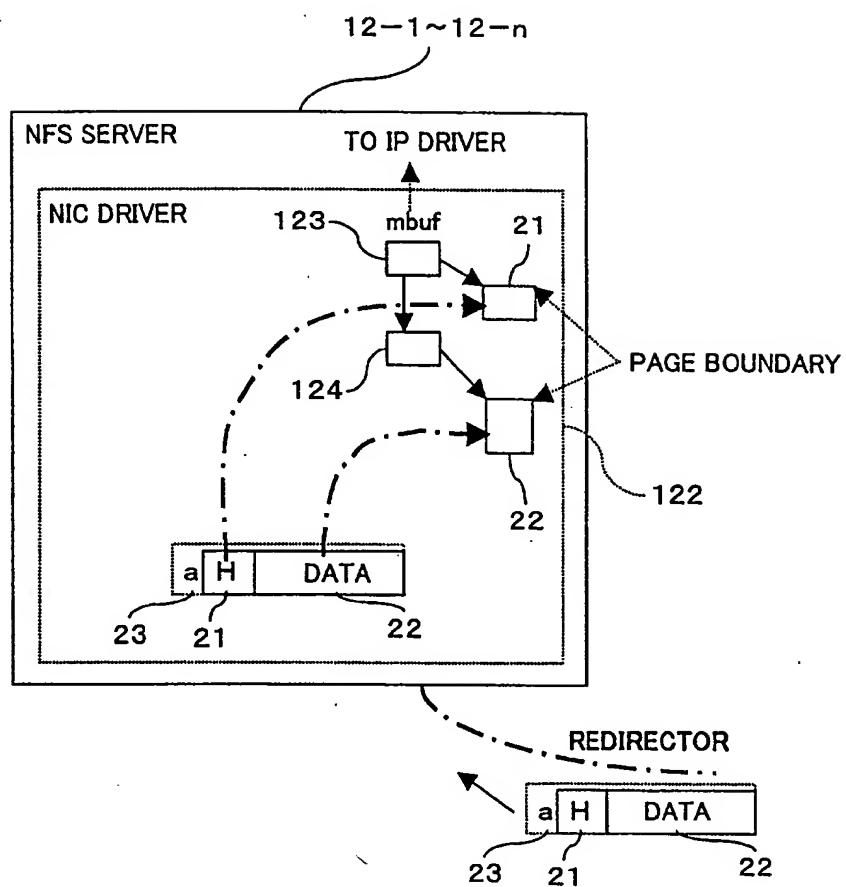


FIG. 11

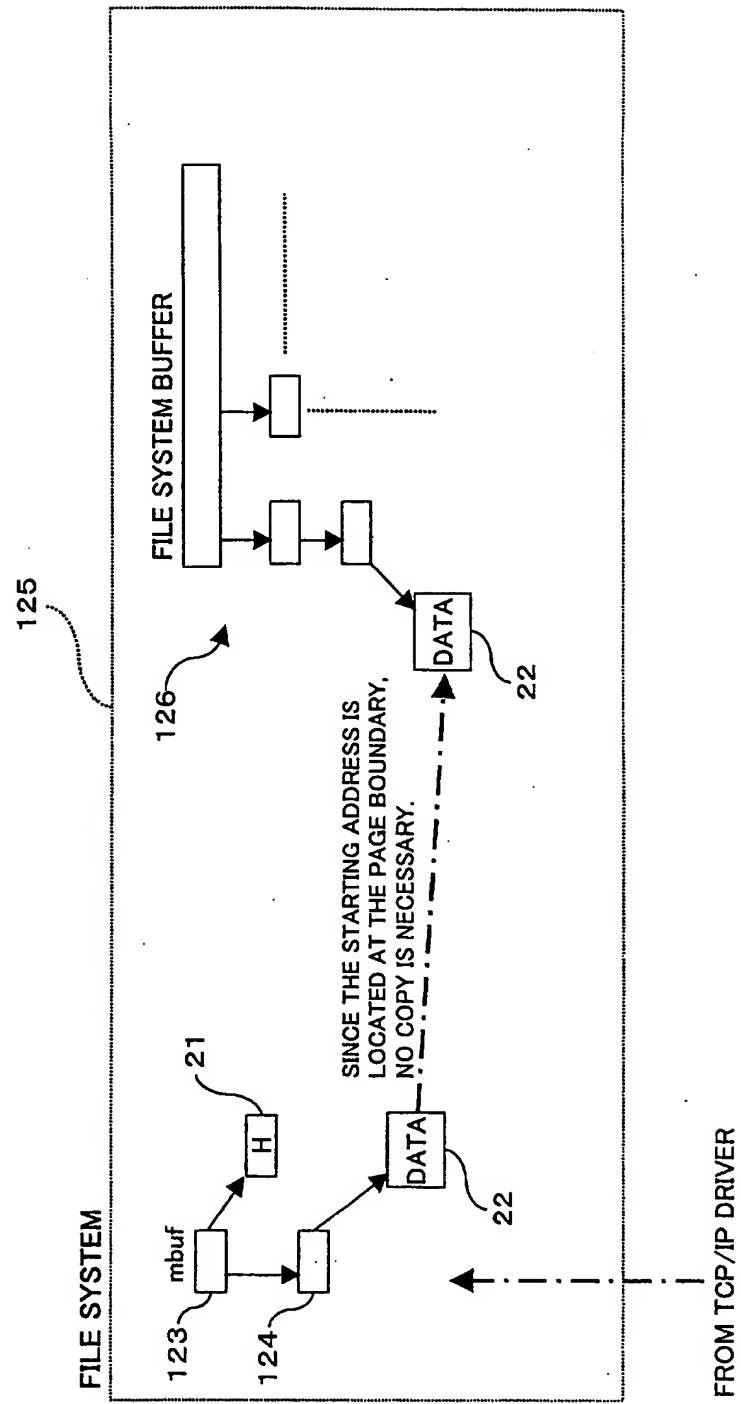


FIG. 12

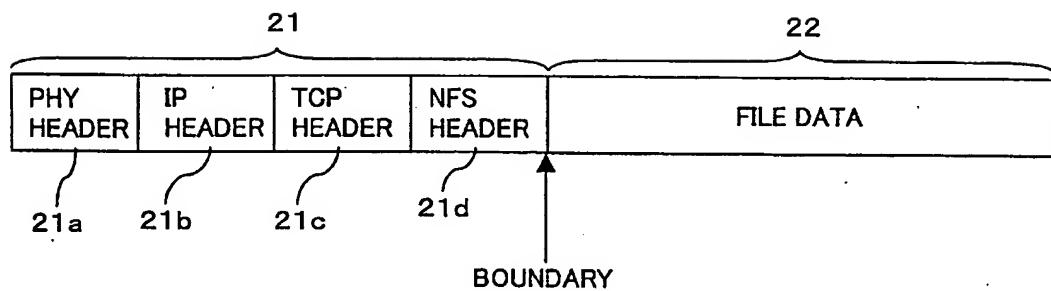


FIG. 13

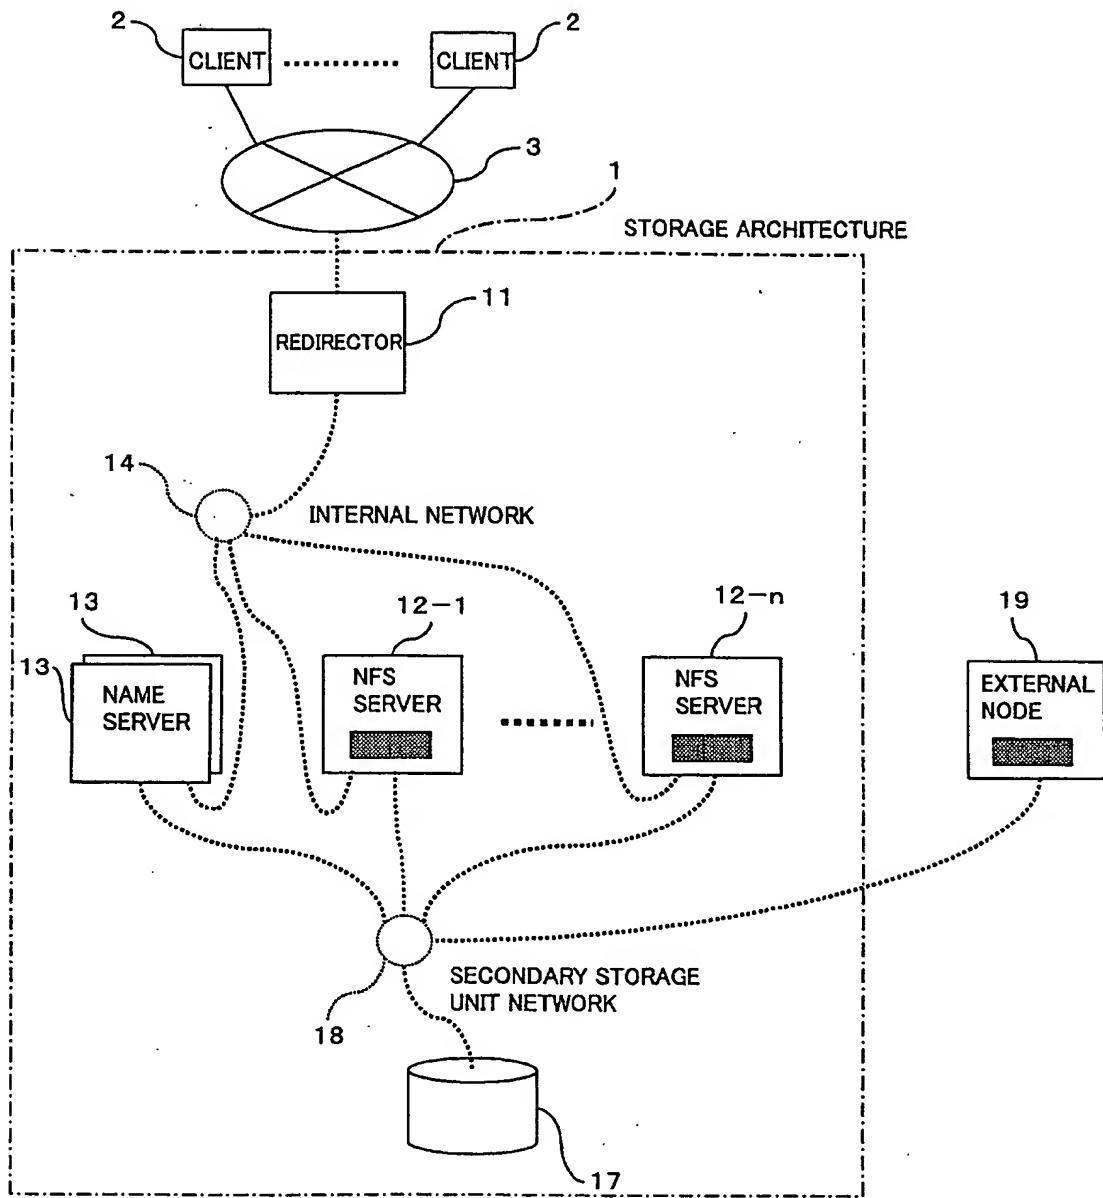


FIG. 14

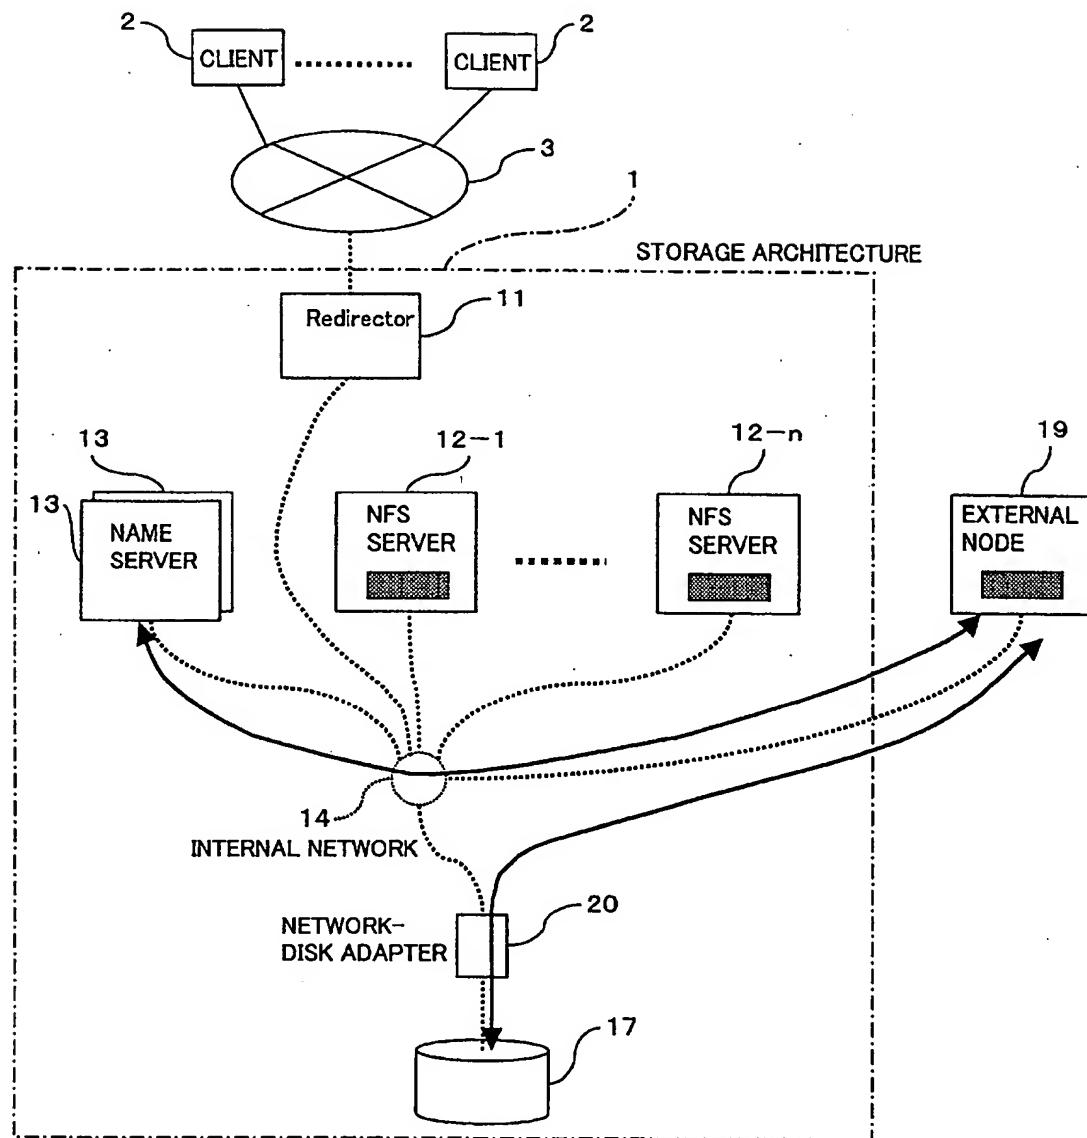


FIG. 15

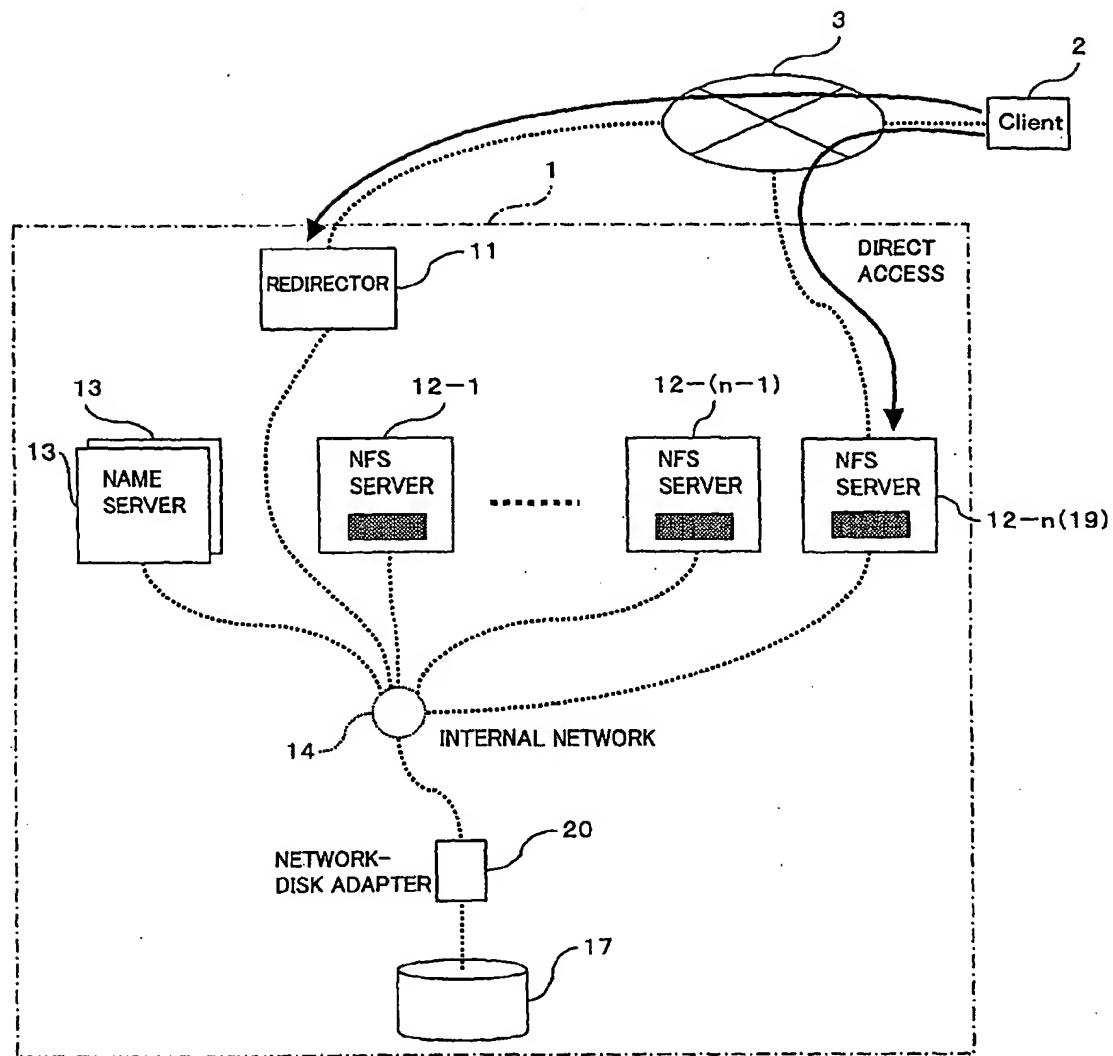


FIG. 16

